

Energy Efficient High Performance Computing Working Group
12/10/13 Meeting Report

INTRODUCTION

The EE HPC WG held a meeting on 12/10/13. This Working Group is composed of members representing major Federal departments and independent agencies, private sector representatives, and members of the academic community. More information can be found at the working group's website, <http://eehpcwg.lbl.gov>.

NEXT MEETING: February 11th, 9:00-10:00AM Pacific Time

***Upcoming WEBINAR: "HPC and the Grid" ~ Wednesday, January 22nd
9:00-10:00AM Pacific Time***

Introductions and Announcements: *Natalie Bates & Dale Sartor, LBNL*

- Membership in the EE HPC WG is up to 430 members.

Conferences Sub-group Update: *Anna Maria Bailey & Marriann Silveira, LLNL*

- **News on EE HPC WG participation in upcoming Conferences**

SC13: Overall, the EE HPC WG's participation in the Supercomputing Conference has been increasing in its impact year-to-year.

The workshop this year was increased to one and a half days from its prior 1 day format and yet the content was rich and the audience participation was strong. It is estimated that there were at least 150 people who attended at least part of the workshop and the room always had at least 75 participants. The workshop presentations are posted on the EE HPC WG website <http://eehpcwg.lbl.gov/documents/sc13-technical-program/workshop>. They have also been submitted to the SC13 conference committee for inclusion in the conference proceedings.

In addition to having sessions from each of the EE HPC WG active Teams, we had Chris Malone from Google and Dan Reed, currently VP of Research and Economic Development at the University of Iowa and formerly with Microsoft, talk about HPC and Warehouse Computers.

Monday's opening remarks were given by Herbert Huber from LRZ and the sessions featured speakers on Architecture and Benchmarking with John Shalf from LBNL, Erich Strohmaier also from LBNL, Wu Feng from Virginia Tech, Steve Poole from ORNL and Jack Dongarra from the University of Tennessee. Steve and Jack both talked about new benchmark workloads that would supplement and/or replace HPL.

The EE HPC WG organized three BoFs this year, one was jointly organized with the Green500 and Top500. Participation in these three BoFs was also very strong, with at least 60 people at the Liquid Cooling Commissioning and TUE BoFs. The Green500 BoF attracted more people, easily over 100, although the numbers were dwarfed by the size of the BoF room.

In addition to the workshop and BoFs, the EE HPC WG had a research booth on the exhibition floor this year. The booth was initiated by Gary Seifert from OSISoft – who did a great job with taking the initiative as well as helping with the booth displays and give-aways. The membership is up by ~50 members and many of them came from interested people who learned about the EE HPC WG from the booth.

****The EE HPC WG website Links and Events page lists many upcoming Conferences and Workshops that have an HPC Energy Efficiency Focus****

Future Conferences: (more details at <http://eehpcwg.lbl.gov/events-and-links>)

- We are waiting to hear from the ACM/SPEC Conference as to whether or not the Power Methodology paper will be accepted. The author notification date is December 15th.
- The ASHRAE T9.9 Committee will be meeting during the ASHRAE annual conference from January 18-22. We will use that meeting as part of the hand-off to ASHRAE of the Liquid Cooling Commissioning document.
- Believe it or not, the SC14 workshop submission deadline is February 7th.
- We are also planning on submitting at least one paper to ISC'14. The submission due date is February 9th. The papers would be on HPC and the Grid as well as TUE.

Infrastructure Sub-Group Update: *William Tschudi, LBNL & Dave Martinez, SNL*

LIQUID COOLED COMMISSIONING TEAM UPDATE: The Liquid Cooling Commissioning Team has collected and reviewed best practices and lessons learned for commissioning of liquid cooling infrastructure. The ultimate goal is to improve the commissioning process for delivering a liquid cooling infrastructure that works when the HPC system is installed.

The Team has drafted a document that is intended to provide some general guidance and recommendations for those who are writing commissioning plans. They have five 'lessons learned' documents that help to illustrate the general guidance and recommendations. The goal is to have closer to 10 lessons learned- SO PLEASE LET NATALIE KNOW IF YOU ARE INTERESTED. The 'lessons learned' can be anonymous if necessary.

Members of the team presented at the SC13 Workshop as well as at a BoF. The presenters include Dave Martinez from Sandia, Tom Durbin from NCSA, Detlef Labrenz from LRZ in Munich, Marriann Silveira from LLNL, and Mike Ellsworth from IBM.

Mike is also active with ASHRAE T9.9. The team has been working with ASHRAE and they are interested in using the document we produce as the basis for an ASHRAE publication on liquid cooling commissioning. The next ASHRAE T9.9 Meeting is January 18-22 in NYC.

- **TUE TEAM:** The TUE Team has developed a metric that improves PUE by accounting for infrastructure elements that are a part of the HPC system (like cooling and power distribution). The team has already tested the metric at ORNL and published a paper with the results. (The paper was awarded with “Best Paper” at ISC13.)

There was a TUE presentation at the SC13 Workshop by Mike Patterson and a BoF. Chung Hsing Hsu from ORNL, Anna Maria Bailey from LLNL, Herbert Huber from LRZ and Satoshi Itoh from the Japanese National Institute of Advanced Industrial Science and Technology presented the results of their TUE testing at the BoF.

Please let Natalie know if you are interested in helping test and deploy the TUE metric.

- **ENERGY REUSE EFFECTIVENESS:** The Energy Re-use Effectiveness Team in collaboration with The Green Grid has developed a standard metric for measuring the contribution of re-using heat generated by HPC systems for other useful purposes.

Steve Hammond moderated a panel session at the SC13 workshop and panelists (who have some experience with heat reuse) included Gert Svensson, from KTH in Sweden, Paul Brenner from the University of Notre Dame and Bill Tschudi from LBNL.

Lars Eilersten from the Norwegian University of Science and Technology said they do heat re-use, but he thought everyone was doing it. He was surprised to discover at the workshop that most everyone else is just beginning the journey.

Systems Sub-group Update: *Natalie Bates, LBNL*

SYSTEM WORKLOAD POWER MEASUREMENT METHODOLOGY: The EE HPC WG along with the Green500, Top500 and Green Grid have developed a standard methodology for measuring system power while running a workload. The team developed the standard, refined it through both alpha and beta testing and collaborated with the Green500 List to ensure adoption as the Green500 run rules. They are now developing outreach and other tools for broader adoption of the measurement methodology. The ultimate goal is to have broad use of the highest quality energy and power measurement methodology for all of their system workload energy efficiency benchmarking activities.

There were two sites who made new Level 3 (highest quality) submissions to the November13 Green500 List. These were the Swiss Supercomputing Center (CSCS) and the Indian Center for the Development of Advanced Computing (C-DAC). The systems were Piz Daint, number 4 and PARAM Yuva, number 44.

Thomas Schulthess, the Director for CSCS and Robin Goldstone, LLNL were presenters at the joint Green500/Top500/EEHPC WG BoF. They both made a strong statement of support for the higher quality power measurements. Thomas’s support was also captured in an interview in *The Register*.

“During the session, Thomas Schulthess of CSCS, Piz Daint’s home, was adamant that Level 3 – which takes the Level 2 rules and makes them more stringent – is the only truly legitimate way of measuring a system’s power consumption, although Level 2 is acceptable, as well. Level 1 is his *bête noire*.

"I am a physicist by training," he said, "a professor of physics at ETH, and I have to measure that one number – that's the true number. There are no two different numbers or two different efficiencies in systems."

When Schulthess ran Level 1 testing on Piz Daint, he obtained one number; when he ran Level 3 testing, the score was more accurate but less impressive. The difference was more than enough to make the "true number"-seeking physicist in him uncomfortable, seeing as how the Green500 accepts Level 1–tested submissions.

Running Linpack on Piz Daint and using Level 3–class analysis, Schulthess and his team came up with a score of 3,186 Mflops/W. Running under Level 1 rules, that score jumped to 3,864Mflops/W.

"People said we should have submitted this" higher Level 1 number to Green500, he said, "but it is wrong. It is the wrong number." So Schulthess submitted a Level 3 score even though his competition was allowed to submit Level 1 scores. Piz Daint ended up at number four instead of number two.

"Every center, every system owner and system operator, is responsible to publish the right number," he said. "It's just the way that things are done in science, and I hope the supercomputing community adheres to the same rules. I'm not sure, from the numbers I've seen, whether this is always the case."

Myslewski, Rik. "Green supercomputer benchmarks make boffins see red, check blueprints." *The Register* 21 November 2013.

http://www.theregister.co.uk/2013/11/21/gpu_accelerators_overrun_green500_list_of_energyefficient_hpc_systems/

The Team also made a submission to the ACM/SPEC Conference to be held in Dublin in March. Author notification is December 15th.

HPC AND GRID INTEGRATION: The Demand Response Team is investigating how HPC centers have, can and should engage more actively with the Grid electricity providers. The Team has collected information from 11 US-based SC sites that are on the Top100 list. This includes LLNL, LANL, ORNL, LBNL, ANL, Purdue, SDSC, NCSA, NOAA, Intel and WPAFB. This is an investigative activity with the ultimate goal of educating the HPC DOE Facility and Operations Managers about HPC and grid integration opportunities and challenges.

There was an SC13 workshop panel session on this topic. It was moderated by Anna Maria Bailey with panelist Jim Rogers, ORNL, Josip Loncaric, LANL and Bob Conroy, OsiSoft.

This Team is in the process of writing a paper for submission to ISC14. This paper should be ready for EE HPC WG review by January 20th in order to accommodate the February 9th ISC Submission deadline.

PROCUREMENT CONSIDERATIONS: The RFP Team has a whitepaper that recommends procurement document requirements that target more energy efficient HPC systems. The intention is to raise the bar and extend the requirements with a yearly update of the whitepaper. The 2013 focus is on measurement capabilities.

This Team also had an SC13 workshop panel session, moderated by Steve Martin, Cray with panelists Daniel Hackenberg, University of Dresden, Chung-Hsing Hsu, ORNL and Jim Laros, SNL.

The system integrators and component providers have been invited to respond to the requirements and talk about their roadmaps. We've had webinars from Cray, IBM, Intel and AMD. We are hoping for additional vendors to participate in this series of webinars. The ultimate goal is to have vendors respond to the requirements with cost-effective product features.

PARTICIPANTS INCLUDED

Name	Organization
Abdulla, Ghaleb	LLNL
Ailee, Guy	Intel
Bailey, Anna Maria	LLNL
Bates, Natalie	LBNL
Bonnie, Amanda	LANL
Brantley, Bill	AMD
Cheng, Joe	
Choi, Jee Whan	Georgia Tech
Cocilova, Anita	LLNL
Durbin, Thomas	NCSA
Ellsworth, Mike	IBM
Goodwin, Dave	ASCR/Office of Science
Horn, Shannon	Colorado State University
Johnson, Nick	University of Edinburgh
Kalwani, Sharan	Fermi National Accelerator Lab
Labrenz, Detlef	LRZ
Mahdavi, Rod	LBNL
Martin, Steve	Cray Inc.
Martinez, Dave	SNL
Rajappa, Murali	Intel
Silveira, Marriann	LLNL
Stemple, Walker	Dell
Tomlinson, Bob	LANL
Tshudi, Bill	LBNL
Whiting, Eric	Idaho National Laboratory